

WE CLAIM:

1. A data storage management system comprising:
at least one network-accessible storage device
capable of storing data;

5 a plurality of network-accessible devices configured
to implement storage management processes;

a communication system enabling the storage
management processes to communicate with each other; and

10 wherein the storage management processes comprise
processes for storing data to the at least one network-
accessible device.

15 2. The data storage management system of claim 1
wherein the at least one network-accessible device
capable of storing data comprises a plurality of network-
accessible devices capable of storing data, some of which
are located at distinct network nodes.

20 3. The data storage system of claim 1 wherein the
storage management processes comprise processes for
serving data from the at least one network accessible
storage device.

4. The data storage system of claim 1 wherein the
at least one storage device comprises a RAID storage
system.

25 5. The data storage system of claim 1 wherein the
at least one storage device comprises a computer with
direct attached storage (DAS) selected from the group
consisting of magnetic hard disk, magneto-optical,
optical disk, digital optical tape, holographic storage,
quantum storage, and atomic force probe storage.

30 6. The data storage system of claim 2 wherein the
plurality of storage devices comprises a peer-to-peer

network of storage devices, each storage device having means for communicating state information with other storage devices, at least one storage device comprising means for receiving storage requests from external 5 entities, and at least one storage device comprising means for causing read and write operations to be performed on others of the storage devices.

7. The data storage system of claim 1 wherein the communication system comprises a TCP/IP over Ethernet 10 network.

8. The data storage system of claim 1 wherein the communication system comprises Gigabit Ethernet network.

9. The data storage system of claim 1 wherein the communication system comprises a Fibre Channel fabric.

15 10. The data storage system of claim 1 wherein the communication system comprises a wireless network.

11. The data storage system of claim 2 wherein the processes for storing data comprise processes that implement a RAID-type distribution across the plurality 20 of network-accessible devices.

12. The data storage system of claim 2 wherein the processes for storing data comprise processes that implement an n-dimensional parity scheme across the plurality of network accessible devices.

25 13. The data storage system of claim 12 wherein the processes for storing parity data expand or contract the size of the parity group associated with each data element to whatever extent is desired.]

14. The data storage system of claim 12 wherein the storage management processes further comprise processes for recovery of data when one or more of the network-accessible storage devices is unavailable.

5 15. The data storage system of claim 12 wherein the storage management processes further comprise processes for access to stored data when one or more of the network accessible storage devices are not desirable data sources for reasons including but not limited to efficiency, 10 performance, network congestion, and security.

15 16. The data storage system of claim 1 wherein the plurality of network-accessible devices configured to implement storage management processes further comprise commercial off-the-shelf computer systems implementing a common operating system.

20 17. The data storage system of claim 1 wherein the plurality of network-accessible devices configured to implement storage management processes further comprise commercial off-the-shelf computer systems implementing a heterogeneous set of operating systems.

18. The data storage system of claim 1 wherein the storage management processes comprise processes for implementing greater than two dimensions of parity.

25 19. The data storage system of claim 2 wherein the processes for storing data comprise processes that store parity and/or mirror data across more than one of the plurality of network-accessible storage devices.

20. The data storage system of claim 1 wherein the storage management processes comprise processes for

adding and removing additional storage capacity to individual storage devices and the system as a whole.

21. A method of data storage management comprising the acts of:

5 providing at least one network-accessible storage device capable of storing data;

implementing a plurality of storage management process instances;

10 communicating storage messages between the storage management process instances; and

storing data to the at least one network-accessible device under control of at least one instance of the storage management processes.

22. The method of claim 21 wherein the at least one network-accessible device capable of storing data comprises a plurality of network-accessible storage devices capable of storing data, some of which are located at distinct network nodes.

23. The method of claim 21 further comprising serving data from the at least one network accessible storage device.

24. The method of claim 21 wherein the step of storing data to the at least one storage device comprises storing the data in a RAID-like fashion.

25. The method of claim 22 further comprising implementing a peer-to-peer network between the plurality of storage devices; and

communicating state information between the plurality of storage devices; and

30 performing read and write operations using the plurality of storage devices.

26. The method of claim 22 wherein the step of storing data comprises storing data using a RAID-type distribution across the plurality of network-accessible storage devices.

5 27. The method of claim 22 wherein the act of storing data comprises storing parity and/or mirror data across more than one of the plurality of network-accessible storage devices.

10 28. The method of claim 22 wherein the storage management process instances further comprise processes for recovery of data when one or more of the network-accessible storage devices is unavailable.

29. A data storage management system comprising:
a plurality of network-accessible storage devices capable of storing data;
a plurality of network-accessible devices configured
5 to implement storage management processes;
a communication system enabling the storage management processes to communicate with each other;
wherein the storage management processes comprise processes for storing data to the at least one network-accessible storage device; and
10 wherein the at least one network-accessible device capable of storing data comprises a parity record holding parity information for at least one other storage node.

30. The data storage system of claim 29 wherein the parity record comprises data capable of correcting errors on another network-accessible storage device.

31. The data storage system of claim 29 wherein the parity record is stored in data structures on at least two network-accessible storage devices.

32. The data storage system of claim 29 wherein the data storage system comprises data structures implementing parity with one or more other, external data storage systems.

33. A method of data storage management comprising the acts of:

providing a plurality of network-accessible storage devices each capable of storing data;

5 implementing a plurality of storage management process instances;

communicating storage messages between the storage management process instances; and

10 identifying one or more storage devices associated with the data to be stored;

determining parity information for the data to be stored; and

storing the unit of data and/or parity data across the two or more storage devices.

34. The method of claim 33 wherein the parity data comprises an error checking and correcting code.

35. The method of claim 33 wherein the parity data comprises a mirror copy of the unit of data to be stored.

36. The method of claim 33 wherein the parity data is stored in a single network storage node and the unit of data is stored in two or more network storage nodes.

37. The method of claim 33 wherein the parity data is distributed across multiple storage nodes.

38. The method of claim 33 further comprising:
retrieving the stored unit of data;

verifying the correctness of the stored unit of data using the parity data;

5 upon detection of an error in the retrieved unit of data, retrieving the correct unit of data using the parity data.

39. The method of claim 33 further comprising:
attempting to retrieving the stored unit of data;
detecting unavailability of one of the two or more network storage nodes; and

5 in response to detecting unavailability, reconstructing the correct unit of data using the parity data.

40. The system of claim 33 wherein the act of storing the unit of data comprises distributing non-
10 identical but logically equivalent data in a storage node.

41. The system of claim 33 further comprises storing lossy equivalent data in a storage node.

42. A method of data storage management comprising the acts of :

5 providing a plurality of network accessible storage devices capable of storing data;

implementing a plurality of storage management process instances;

communicating storage messages between the plurality
10 of storage management processes;

storing data to the plurality of network accessible storage devices under control of the plurality of storage management processes; and

15 adding and subtracting data storage capacity to and from the data storage under control of the plurality of

storage management processes without affecting accessibility of the data storage.

43. The method of claim 42 further comprising:
monitoring the data storage for faults by means of
5 the plurality of storage management processes;
compensating for the faults by manipulating the data storage under control of the plurality of storage management processes without affecting accessibility of the data storage.

44. A method of data storage management comprising the acts of:

providing a plurality of network-accessible storage devices each capable of storing data;

5 implementing a plurality of storage management process instances; and

communicating storage messages between the storage management process instances, wherein any of the storage management process instances is capable of storage 10 allocation and deallocation across the plurality of storage nodes;

45. The method of claim 44 wherein the storage allocation management processes are configured to use the storage messages to reconstruct data stored in a failed one of the storage devices.

46. The method of claim 44 wherein the storage management processes are configured to migrate data amongst the storage devices using the storage messages in response to a detected fault condition in at least one of the storage devices.

47. The method of claim 44 wherein the storage management processes are configured to migrate data

amongst the storage devices using the storage messages in preemptively when a fault condition in at least one of the storage devices is determined to be likely.

48. The method of claim 44 wherein the plurality of 5 storage devices comprises an arbitrarily large number of storage devices.

50. The method of claim 44 further comprising:
associating parity information with a data set;
storing the parity information in at least some of 10 the storage devices; and
serving data requests corresponding to the data set by accessing the parity information associated with the data set.

51. The method of claim 44 further comprising:
15 storing a data set in a plurality of the data storage devices using the storage management processes;
serving data requests corresponding to the data set by accessing the plurality of data storage devices in parallel.

20 52. The method of claim 44 further comprising
encrypting storage messages before communicating.